

The UCLA Earth System Model: Development and Applications

Carlos R. Mechoso and John D. Farrara

Department of Atmospheric Sciences, University of California, Los Angeles
405 Hilgard Ave., Los Angeles, CA 90095-1565

Abstract- An Earth System Model (ESM) has been developed under support from the first two rounds of NASA's ESS HPCC program. In the first round the UCLA atmospheric general circulation model (AGCM) was coupled to a JPL version of the Parallel Ocean Program (POP), and to the UCLA atmospheric chemistry and transport model (ACTM). In the second round a simplified version of the JPL ocean biogeochemistry model (OBM) was added and a Distributed Data Broker (DDB) was incorporated. The DDB is a unique tool for moving data between model components in multi-processor environments, data archives, and visualization clients. The code of the expanded ESM was parallelized and highly optimized. The methodology and lessons learned in the code parallelization, optimization and DDB design are described and selected results of research on outstanding aspects of the climate system using versions of the ESM are presented.

Plans for Round 3 include a demonstration of the functionality of the ESM Framework (ESMF) being developed under ESTO through an analysis of the El Niño prediction capability of the atmosphere-ocean component of our ESM in combination with NASA/JPL ocean data and optimization products.

I. INTRODUCTION

This paper reports our efforts to develop a comprehensive Earth System Model (ESM) and its embedding within an advanced computational infrastructure. We envision a system to be used in support of research, teaching at various levels, and informed decision-making on environmental issues. Our current objective is a model suitable for investigation of global and regional climate phenomena with up to centennial time scales. The ESM code will be executed in a computational framework being developed under developed by NASA's Earth Science Technology Office (ESTO) Computational Technologies (CT) Project. This ESM framework (hereafter ESMF) aims to minimize current roadblocks to component interoperability, and thus to facilitate model improvement, validation, and application.

II. THE UCLA ESM

Our current ESM core comprises three models representing the coupled dynamics, physics and chemistry of the global atmosphere and world oceans (see Fig. 1): 1) an atmospheric general circulation model (UCLA AGCM), 2) an oceanic general circulation model (Parallel Ocean Model: POP), and 3) an atmospheric chemistry and transport model (UCLA

ACTM). Data exchanges are through a novel Distributed Data Broker (DDB) which is described in further detail below.

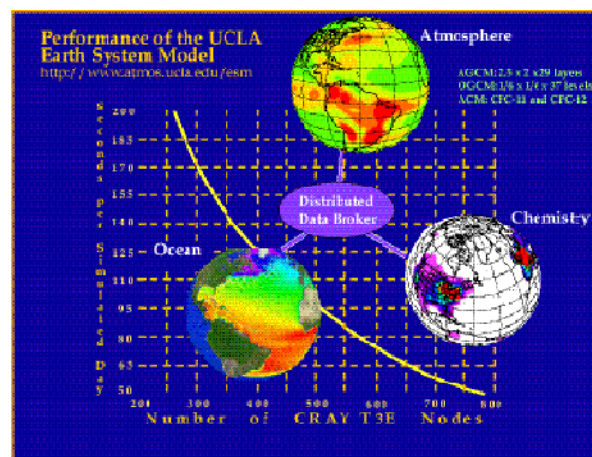


Fig. 1. Schematic of our current Earth System Model. The yellow line is a scaling curve of model performance on the Round 2 testbed machine (CRAY T3E-600).

Figure 1 is a schematic of our ESM and its current performance. The baseline ESM code achieved 40 GFLOPS on 768 nodes of the Round 2 ESS testbed (CRAY T3E-600). Realizing this level of performance required the application of several different optimization strategies. These included static and dynamic load balancing schemes for the AGCM [1] and single-node optimization techniques for both the AGCM and POP. The single-node optimization was based on code restructuring to improve cache re-use, loop unrolling, and selected use of single precision arithmetic. In addition use of a more efficient algorithm in one of the major AGCM parameterizations (cumulus convection) resulted in a significant reduction in the wall-clock time to solution [2]. Lastly, the CRAY T3E hardware STREAMS were enabled.

A. Load Balancing Schemes

Load imbalances generated by computations related to the physical parameterizations that depend on time and space and by filtering operations largely determine the performance of an AGCM code. For the most part, these imbalances can be dealt with via load redistribution techniques that target the reduction of memory latency (local or remote to a node) while maximizing the use of available resources in the system.




Fig. 2. An example of AGCM/Physics load balancing involving four processors. Boldface numbers represent AGCM/Physics loads in seconds per simulated day of climate, italics numbers in the upper left corner of boxes represent processor number. Arrows represent data exchanges between processors.

Variations in time and space of the intensity of atmospheric convection and differences in extent between sunlit regions and those in darkness are primarily responsible for the load imbalances in the Physics. We have addressed this load imbalance with a dynamic scheme [1]. This is based on pairwise exchanges of data between processors with very different predicted loads (see Fig. 2). The load predictions are obtained by periodically measuring the unbalanced load in each processor. As shown in the example in Fig. 2, it may be necessary to perform the pairwise data exchange more than once to obtain a satisfactory final load distribution.

The smoothing of fields in high-latitudes is performed via a Fourier filter, resulting in the primary load imbalance. The characteristics of this particular load imbalance do not change with time and are addressed by a static scheme [1]. This is based on evenly distributing latitude bands of fields (covering all heights and a subset of longitudes) to be filtered among all processors. Such a distribution is made in four steps, of which the first two are schematically shown in Figure 3. The first step consists of a rearrangement in the latitudinal direction, and results in a more even distribution of filtering load in latitude. The second step consists of a rearrangement in the longitudinal direction, and results in each processor containing fields that are complete in longitude and ready to be filtered. The third step is an inversion of the second, and the

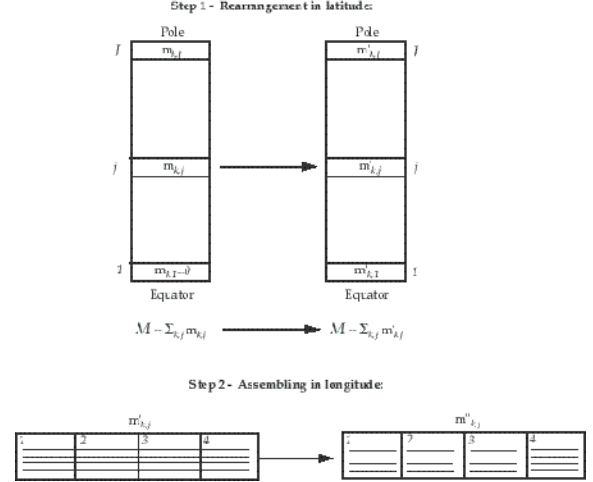


Fig. 3. Schematic of the first two steps in the load balancing for AGCM/Dynamics of Lou and Farrara [1]. $m_{k,j}$ is the number of latitudes to be filtered of variable k in processor j . The solid and dashed lines in Step 2 represent complete longitude-height slices of the data. Prior to Step 2 these slices are distributed among all four processors; upon completion of Step 2, each resides in a single processor. Steps 3 and 4 (not shown) are inversions of Steps 2 and 1, respectively.

fourth step an inversion of the first. Since the number of such latitude-height cross sections is necessarily limited whatever the model's resolution, the load redistribution by this scheme becomes poor when large numbers of processors are used. Therefore, we also allow for the breakup of latitude-height cross sections into individual latitudes of data. This results in a larger number of units of work to be redistributed, giving a better load balance on large numbers of processors.

The primary source of load imbalance in the OGCM/POP is the differing number of land points included in processor when the simplest longitude-latitude domain decomposition is used. This static load imbalance can be nearly eliminated by application of the re-partitioning scheme of [3], which removes most of the land points from the computational domain. For a North Atlantic version of POP use of this scheme resulted in a savings in total execution of time of 33%.

B. The Distributed Data Broker

For running in multi-processor environments we have designed a DDB to move data between model components, data archives, and visualization clients. The DDB is a general purpose tool for coupling multiple, possibly heterogenous, parallel models. It is implemented as a library used by all participating elements, one of which serves as a distinguished process during a startup phase preceding the main computation. This "registration broker" process corollates offers to produce quantities with requests to consume them, forwards the list of intersections to each of the

producers, and informs each consumer of how many pieces to expect. After the initial phase, the registration broker may participate as a regular member of the computation. A library of data translation routines is included in the DDB to support exchanges of data between models using different computational grids. Having each producer send directly to each consumer conserves bandwidth, reduces memory requirements, and minimizes the delay that would otherwise occur if a centralized element were to reassemble each of the fields and retransmit them. A detailed description of the DDB can be found in [4].

III. CLIMATE SIMULATIONS

The coupled atmosphere-ocean (CGCM) component of the ESM produces a very realistic simulation of El Niño/Southern Oscillation (ENSO) [5], a success shared by only a handful of current models. This success has allowed us to use the CGCM to perform studies aimed at testing simple models of ENSO [6] and the decadal variability of the phenomenon [7].

Another notable success has been the application of the model to studies of the long-term evolution of chlorofluorocarbons (CFCs) in the atmosphere [8]. Multi-decadal simulations were performed to explore the entire global history of the two most widely used CFCs, CFC-11 and CFC-12. The simulations reveal how CFCs emitted at the Earth's surface are transported into the stratosphere, where they are turned into ozone-destroyers by photochemical reactions. Figure 4 shows the distribution of CFC-11 in the lowest model layer for the winter of 1990. By simulating the full stratosphere, the lifetimes of CFCs were determined to a greater precision. These findings have implications not only for ozone depletion but also for global

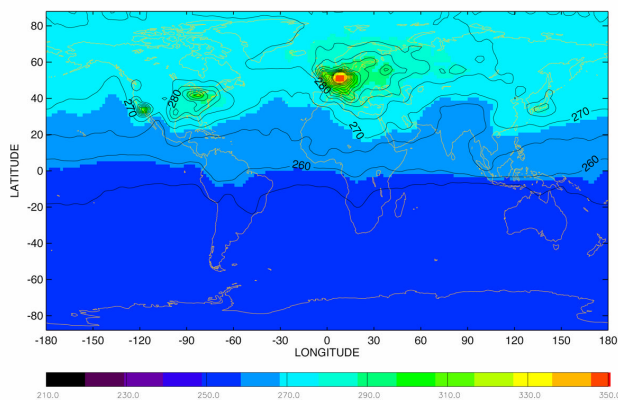


Fig. 4. Simulated concentration of CFC-11 (parts per trillion) in the lowest atmospheric model layer for the northern hemisphere winter of 1990. Emission rates used are those reported from developed countries and projected for non-reporting countries. warming because of the greenhouse gas properties of CFCs.

IV. INTERFACE WITH THE ESMF

To interface with the ESMF, our ESM will be viewed as consisting of four distinct levels of application objects, in which increasing level depth means increased detail.

Level 1 (Top): The objects at this level are the component models that compose the ESM. Examples are the AGCM, OGCM, and ACTM.

Level 2 (Intermediate): The objects at this level are the ESM modules. Examples are PBL/clouds, Land Surface Scheme, Sea-Ice Model.

Level 3 (Low): The objects at this level are the computational components of the models. Examples are horizontal differencing, vertical differencing, time differencing, kernel calculation in the cumulus parameterization.

Level 4 (Deep): The objects at this level are the computational subfunctions used in the model codes. Examples are Fast Fourier Transform, astronomy calculations and generalized table-lookup functions.

V. DEMONSTRATION OF ESMF FUNCTIONALITY

Although there has been enormous progress in our understanding of El Niño and related phenomena, much remains to be learned as is evident, for example, from the current controversies concerning possible changes in the properties of El Niño. The alternation between complementary El Niño and La Niña states is highly irregular. Up to the 1960's El Niño occurred sporadically, but then became more regular with a period of approximately 3 years. Since the 1980's the period seems to have increased to approximately 5 years, but the cold La Niña episodes have practically disappeared. What causes the Southern Oscillation to be so irregular? Some investigators contend that the time-series is stationary and reflects the impact of random disturbances, mainly of atmospheric origin, on a regular oscillation attributable to ocean-atmosphere interactions. Other investigators argue that global warming is the reason why El Niño was so exceptionally intense in 1982 and 1997, and why La Niña was practically absent during the 1980's and 1990's. This is clearly an important issue, given the current concern about global warming.

The debate about the appropriate interpretation for the irregularity of El Niño/La Niña concerned strictly statistical matters until [9] discussed the problem in the context of a stability diagram for ocean-atmosphere interactions. They calculated how the properties of El Niño would change should there be changes in background parameters such as the spatially averaged depth of the oceanic thermocline and the time-averaged intensity of the trade winds. They concluded that, because the background state is subject to a continual decadal oscillation, which has been associated with a deeper thermocline and weaker trade winds since the late 1970's, the frequency of occurrence of El Niño has increased over the past few decades. The findings of [9]

were obtained with a very simple coupled model—basically that of Cane and Zebiak -- which has a large number of parameters that are assigned values on a somewhat arbitrary basis. The advantage of the model is that a broad range of parameter values can be explored easily. The disadvantage is that the model has several unrealistic features (because it has only two layers in its oceanic component for example).

To address the problem with the CGCM, the thermal structure of the ocean - the depth of the thermocline for example - will be changed. In reality, the depth of the thermocline depends on the density of the deep ocean, which in turn depends on the thermohaline circulation. The time-scale for change in the deep ocean is on the order of a thousand years. Hence, for CGCM experiments that cover a few decades, the thermohaline circulation can be regarded as a given. Specification of the density of the deep ocean in effect specifies certain aspects of the thermohaline circulation so that a simulation that covers a few decades permits only the upper ocean to adjust to changing atmospheric conditions. Hence, we propose to explore different background states by specifying different initial thermal structures for the ocean. In the calculations, the upper ocean, in and above the thermocline, will change but not the deep ocean, at least on the time scales under consideration.

We plan to demonstrate the functionality of the ESMF by analyzing the El Niño prediction capability of the CGCM component of the ESM in combination with NASA/JPL ocean data and optimization products. A series of prediction experiments initialized every month and lasting for two years each will be carried out. In a first set of experiments, the ocean component of the ESM will be initialized from temperature and salinity fields derived from a prognostic integration of the MIT OGCM carried out at JPL. In a second set of experiments, the ocean will be initialized from fields derived from circulation estimates that have been constrained with TOPEX/POSEIDON, TAO, XBT, and other data (descriptions of the prognostic and data-constrained estimates and access to the complete results are available at <http://eyre.jpl.nasa.gov/las/main.pl>). Results from these experiments will be analyzed for predictive skill.

VI. SUMMARY

ESMs are fundamental tools for climate studies. We believe that their role will become even more important as hardware and software advances make them faster and friendlier to the users, and as demands for more accurate and detailed predictions grow. The models are making rapid progress towards a successful simulation of the coupled climate system. Such progress requires the continued collaboration of earth scientists and computer scientists, such as that facilitated by NASA's Earth Science Technology Office (ESTO) Computational

Technologies (CT) Project, in developing improved ESMs.

VII. PROJECT PARTICIPANTS

The other participants in this project are: Dr. John Baumgardner (Los Alamos National Laboratory), Prof. Richard R. Muntz (UCLA), Dr. Dimitris Menemenlis (NASA/Jet Propulsion Laboratory), Prof. George Philander (Princeton University).

ACKNOWLEDGMENT

Financial support for this work was provided by NASA under CAN NCCS5-149.

REFERENCES

- [1] Lou, J. Z. and J. D. Farrara, 1998: Performance analysis and optimization on the UCLA parallel atmospheric general circulation model code. *Concurrency: Practice and Experience*, **10**(7), 549-565.
- [2] Mechoso, C. R., J. D. Farrara, L. A. Drummond, J. Spahr and J.-Y. Yu, 2000: An atmosphere-ocean model: Code optimization and application to El Niño. Development and Application of Computer Techniques to Environmental Studies VIII. Gabriel Ibarra-Berastegi, C. A. Brebbia and P. Zanetti, Eds., WIT Press, Southampton, Boston, 261-278.
- [3] Wang, P., D. S. Katz & Y. Chao, 1997: Optimization of a parallel ocean general circulation model. *Proceedings, Supercomputing '97*, November 1997, San Jose, CA, IEEE Computer Society.
- [4] Sklower, K., H. Robinson, C. R. Mechoso, L. A. Drummond, J. A. Spahr and J. D. Farrara, 2002: The Distributed Data Broker: A decentralized mechanism for periodic exchange of fields between multiple ensembles of parallel computations. Technical Report, Computer Science Dept, University of California, Berkeley, 45 pp.
- [5] Mechoso, C. R., J.-Y. Yu and A. Arakawa, 2000: A coupled GCM pilgrimage: From climate catastrophe to ENSO simulations. *General Circulation Model Development: Past, present and future*. Proceedings of a Symposium in Honor of Professor Akio Arakawa. D. A. Randall, Ed., Academic Press, 539-575.
- [6] Mechoso, C. R., J. D. Neelin, and J.-Y. Yu, 2002: Testing simple models of ENSO, *J. Atmos. Sci.*, submitted.
- [7] Yu, J.-Y., C. R. Mechoso, J. C. McWilliams, and A. Arakawa, 2002: Impacts of the Indian Ocean on the ENSO cycle. *Geophys. Res. Lett.*, in press.
- [8] Gupta, M. L., R. P. Turco, C. R. Mechoso and J. A. Spahr, 2001: On-line simulations of passive chemical tracers in the University of California, Los Angeles, atmospheric general circulation model: 1. CFC-11 and CFC-12. *J. Geophys. Res.*, **106**, 12,401-12,417.
- [9] Federov, A. V. and S. G. Philander, 2000: Is El Niño changing? *Science*, **288**, 1997-2001.

